EPA SERVER V4 DISCUSSION GUIDE – TGG RESPONSE PART I

April 29, 2022

ITI's The Green Grid

1 INTRODUCTION

TGG appreciates US EPA's continuing efforts to improve the ENERGY STAR Computer Server specification and the launch of the discussion guide for Version 4.0, providing the initial list of priority areas for this specification revision with key questions for stakeholders. Our comments touch on all aspects identified in the discussion guide (Updated and new definitions; Potential treatment of Storage "heavy" Servers; Revisions to the SERT tool and Internal Power Supply Efficiency) and more.

To help correlate EPA questions and TGG comments, TGG has added "**EPA Qx**", in bold blue text throughout the document when appropriate.

2 EXECUTIVE SUMMARY

This document provides recommendations for four updated or new definitions that need to stay current with the market and technology dynamics:

- Resilient Server
- High-Performance Computing System
- Storage Heavy Server
- Hyperconverged Server.

With respect to storage heavy servers (SHS), TTG recommends creating SHS categories based on the maximum number of devices and set an energy efficiency limit for each of those categories. It is impossible to test all permutations, so TGG recommends testing two configurations (one with the minimum storage device count and one with 50% of the maximum storage device count).

TGG is not expecting significant updates to the SERT 2.x test suite for the EPA Server Energy Regulation Version 4.0. However, even with minor updates it is important to define a process that allows those without creating unnecessary EPA effort and churn. As such, TGG recommends utilizing the ISO/IEC 21836:2020 Section 9 compliance methodology.

When it comes to internal power supply efficiency, to avoid unnecessary energy consumption and cost, TGG recommends that 80 Plus Titanium level is adopted only for power supplies that are rated for 750 watts or more, and Platinum for power supplies less than 750 watts.

TGG also recommends the use of external Power Supplies that meet DOE level VI as an option for low end ENERGY STAR compliant servers.

Finally, the current EPA Energy Star Computer Server specification V3.0 includes conflicting text between family definitions and test requirements when dealing with Servers with one processor populated in a two Socket system. TGG proposes to modify Section 6.1.2 of the specification to specify how to handle 2 Socket servers which ship with CPU models which only support 1 socket operation.

3 COMMENTS AND SUGGESTED CHANGES

3.1 Definitions (EPA Q1, EPA Q2)

Resilient Servers

TGG recommends the following changes (bold indicates added words) to section B(2) and B(4) of the Resilient Server definition. Section A, numerals 1., 3., and 5. to 10. of Section "B" and Sections "C" and "D" remain without change.

> "Resilient Server: A computer server designed with extensive Reliability, Availability, Serviceability (RAS) and scalability features integrated in the micro architecture of the system, CPU and chipset. For purposes of ENERGY STAR certification under this specification, a Resilient Server shall have the following characteristics …"
>
> B. System Recovery & Resiliency: No fewer than six of the following characteristics shall be present in the server:
>
> 2. "The processor technology is designed to provide additional capability and functionality without additional chipsets, **enabling it** to be designed into systems with four or more processor sockets;"
>
> **Rationale:** Changing from "enabling the design into systems…" to "enabling it" as this text is intended to refer to the processor technology itself.
>
> 4. "Memory Sparing: A portion of available memory may be pre-allocated **or re-purposed** to a spare function such that data may be migrated to the spare upon a perceived impending failure;"
>
> **Rationale:** Adding "or re-purposed" to the memory sparing definition enables more inclusiveness of solutions that may not require extra memory reserved for sparing purposes.

Storage Heavy Server

TGG recommends the following definition for a Storage Heavy Server. A storage heavy server has additional support infrastructure which increases the power consumption of the server even without the additional storage capacity installed. As SERT is normally run with two storage devices installed, this will decrease the SERT score due to the additional power consumption of the additional infrastructure. This requires a different treatment for this type of servers which we also propose in section 3.2 of this document.

> **storage heavy server**
> a Storage Heavy Server is a Server with greater storage capacity than a standard server. As shipped, the server supports 30 or greater internal storage devices.

In support of the previous definition, the ENERGY STAR Computer Server specification v4 needs to also include a definition of Internal Storage Device, which we provide below for your consideration.

internal storage device
A non-volatile device installed inside the primary server chassis which is intended for data storage.

Note 1 to entry: Examples of internal storage devices include magnetic disk drives, solid state drives and NVM-e devices.

Per the EPA Discussion Guide released, there is also a need to define the term "Hyperconverged Server" and modify and/or simplify the definition of "High-Performance Computing (HPC) System" to align with recent technology changes. We recommend the following definitions

high-performance computing system
HPC system
HPC server
computing system which is designed, marketed, sold, and optimized to execute highly parallel applications for high performance, deep learning, or artificial intelligence applications

Note 1 to entry: HPC systems consist of multiple clustered servers, primarily for increased computational capability, high speed inter-processing interconnects, large and high bandwidth memory capability and often accelerators such as GPGPUs or FPGAs.

Note 2 to entry: HPC systems may be purposely built or assembled from more commonly available computer servers.

hyperconverged server
A highly integrated enterprise device which contains the same components as a computer server in addition to the features of a network server and storage server.

3.1.1 Treatment of Storage Heavy Servers

Per EPA Q5, EPA is requesting additional SERT data for this type of servers to support the criteria development process for separate active efficiency requirements. TGG is collecting these data and will be conducting analysis which will be happy to share in the future with the US EPA, but unfortunately TGG will not be able to do this by the specified date of May 6th.

3.1.2 What are Storage Heavy Servers used for?

Storage heavy servers are typically use as cache to an enterprise storage system. Example of applications include intelligent video analytics and media streaming.

3.1.3 Storage Heavy Server SERT Threshold and Categorization Methodology

3.1.3.1 Storage heavy server categories:

Per EPA Q3 and Q4, TGG is collecting data to determine its suggested approach, and does have an additional proposal to the EPA's recommended solutions under consideration.

Our initial recommendation is for storage heavy servers to be categorized by the maximum number of supported devices. Categories with corresponding efficiency limits will be

determined once the analysis mentioned above is finalized. A combination of measured and modeled data is being used for this purpose.

3.1.4 Storage Heavy Storage Test Configuration Definitions

As the EPA is aware, testing many of the enormous number of server configuration permutations is not feasible, and for storage heavy servers, the cost to populate the maximum number of supported storage devices is substantial. Thus, similar to the volume server family configurations, TGG proposes that only two configurations be tested, and similar to the current Blade server test methodology, a "minimum" configuration and a configuration with "50% of storage devices supported", be tested. Additionally, since lower-end CPUs obtain lower SERT scores, they are a good proxy for all CPUs, TGG recommends that the CPU and memory configuration for both tested configurations use the CPU and memory requirements of the Low-end Performance Configuration.

Family Configurations for testing:

Storage heavy server families should be represented by two tested configurations. SPEC SERT results should be collected using these two configurations.

Test Condition	Storage Device Count	Server Efficiency
Minimum storage device count		
Typical (50% of maximum supported storage devices)		

Minimum Storage Configuration: This configuration shall be the Low-end Performance Configuration as defined in Section 1.G.2.A of the Product Specification for Computer Servers Eligibility Criteria Version 3.0, with the additional requirement that the storage device count shall match the minimum number supported by the storage heavy server.

Typical Storage Configuration: This configuration shall be the Low-end Performance Configuration as defined in Section 1.G.2.A of the Product Specification for Computer Servers Eligibility Criteria Version 3.0, with the additional requirement that the storage device count shall match 50% of the maximum number supported by the storage heavy server.

3.2 Revisions to the SPEC Server Efficiency Rating Tool (SERT)

Per EPA Q6, TGG does not believe there will be any significant updates to the SERT 2.x suite. The only expected updates would be to add support for new servers, power analyzers, and temperature sensors.

There is, however, an important update related to the SERT suite. Section 9 of ISO/IEC 21836:2020 standard provides a robust way to ensure that changes to the five main components related to the SPEC SERT® suite are only updated in a manner which does not adversely affect quality or comparability. These test plans provide functionality and quantitative requirements for updates and ensure thorough testing has been conducted on new server architectures, power analyzers, temperature sensors, and software components.

To utilize this important safeguard, save the EPA the effort necessary to evaluate and communicate acceptance for each SERT suite minor update, and remove the delay between SERT suite minor version release and EPA acceptance, The Green Grid recommends adding the below requirement to the ENERGY STAR Sever v4 specification.

"Only ISO/IEC 21836:2020 compliant SPEC SERT® suite minor versions, SPEC PTDaemon® interface versions, power analyzers, temperature sensors, and Client Configurations as documented in section 9 of ISO/IEC 21836:2020 international standard, shall be used."

3.3 Internal Power Supply Efficiency (EPA Q7, EPA Q8)

Per the EPA Discussion Guide, the EPA is considering raising the bar on the power supply efficiency requirements for Version 4 of the Product Specification for Computer Servers.

In addition to the supply chain problems that are presently experienced by the industry and the impact it reflects on the availability of Titanium power suppliers in the market, the availability of low wattage Titanium power supplies is limited. Currently, only 4 power supplies meeting the Titanium level are listed under the 230V Internal category on the 80 PLUS website that are less than 750 watts. 3 of those are from a single supplier. However, at 750 watts, there are 15 Titanium level power supplies and at 800 watts, there are 20 power supplies listed as Titanium.

Several lower cost servers using lower power processors are on the market and these products can utilize 500 watt power supplies, using these smaller (right sized) power supplies is more efficient. Making smaller Titanium power supplies is harder than large Titanium power supplies because there is a fixed amount of energy consumed for AC to DC and voltage conversion. Small power supplies cannot amortize these fixed costs per Watt. The result is that manufacturers will choose to use higher wattage power supplies that are able to meet Titanium level efficiency more cost effectively.

There is a small efficiency difference between these two efficiency curves and using a larger power supply than necessary will change the load point on the PSU efficiency curve. The result is very likely to be a negligible energy savings for the end user at significant cost increase. The two tables below show example comparisons of a 500W platinum and 1000W titanium power supply for a 250W load point using the efficiency values defined in the appropriate 80 Plus standard for both the 115V internal and 230V internal categories. The energy cost savings are negligible in both cases.

115V	DC Power	AC Power	Annual Energy Cost @ 0.15/kWh
Platinum 500W	250	271.74	357.07
Titanium 1000W	250	270.76	355.78

230V	DC Power	AC Power	Annual Energy Cost @ 0.15/kWh
Platinum 500W	250	265.96	349.47
Titanium 1000W	250	265.02	348.24

If the ENERGY STAR program would decide to require Titanium level for all single output power supplies, Data Center products (servers) will be required to utilize much higher wattage power supplies than necessary in order to meet operational needs specific to the requirement. This is

completely counter to achieving the energy savings hoped for by establishing this efficiency increase in the specification in the first place. Data centers will face increased costs due to requirement for using the 96% efficient (Titanium Certification) power supplies (over 750 watts) when they are not operationally required for the Data Center servers. At the same time, data centers will experience decreased energy savings as power supply wattages will be oversized, larger than the wattage needed, and peak efficiency will be achieved less frequently due to lack of need for more power in the servers.

For full details please see: https://www.thegreengrid.org/en/resources/library-and-tools/538-230V-ENTERPRISE-PSU-MARKET-READY-EFFICIENCY-STUDY

Recommendations:

Create a 2-tier system of efficiency baseline targets as shown below to provide for maximum efficiency for single output power supplies.

TGG recommends the US EPA only use the 80plus specifications, as any deviation from the standard are difficult for industry to ensure and can lead to redesign and retesting of power supplies. If the EPA does desire changes to 80plus power supply efficiencies, TGG would request EPA, TGG and the 80 Plus program (CLEAResult) work together to propose the changes through the standardized process.

If possible, it would be appreciated if the EPA could provide any power supply efficiency improvements planned for past Version 4 of the EPA server specification as soon as possible, to provide manufacturers adequate development time, especially in the current difficult supply chain environment.

Finally, TGG recommends adding the below text to EPA ENERGY STAR Servers 4.0:

a) For power supplies at or over 750 watts require 80 PLUS® Titanium Certification.
b) For power supplies under 750 watts, require 80 PLUS Platinum Certification.

3.4 External Power Supply Option

Add External Power Supplies that meet DOE level VI as an option for low end ENERGY STAR compliant servers.

3.5 Handling of Partially Populated Systems (2 Socket Servers with 1P only CPUs)

TGG requests the US EPA resolves current conflict that exists between family definitions and test requirements.

The “Product Family” definition in Version 3.0 of the Product Specification for Computer Servers states that a product family can be defined for a server with only partially populated processor sockets (e.g., 1 Processor populated in a two-socket system) as long as it is tested as a separate product family. At the same time, Section 6.1.2 of the specification states that Units Under Test (UUTs) must have the maximum number of processor sockets populated in order to test.

Proposed Solution

Section 6.1.2 of the Product Specification for Computer Servers should be modified to allow for testing of partially populated processor socket product families.

A special case exists where a two-socket server is sold with a CPU model which only supports one socket operation (1P only capable CPU). This effectively turn the server into a one-socket server and it is reasonable to take the position that two-socket populated data is not representative of these servers.

For two-socket servers which support both 1P only capable CPUs and 2P capable CPUs, two sets of family configurations are needed. This will create two entries into the QPL database under a single product Model designation. Both will have the same product model name with a product model number that differs by a -1P designation in the Model Number field. Each of these entries will reference the other product model number in the Additional Model Information field. This will be part of a single Energy Star 4.0 product certification and will not require separate certification certificates for the two scenarios.

ENERGY STAR Unique ID	ENERGY STAR Partner	Brand Name	Model Name	Model Number	Additional Model Information
1234567	Partner Name	Brand	Model_Name	Man_Model_Num	Man_Model_Num-1P
1234568	Partner Name	Brand	Model_Name	Man_Model_Num-1P	Man_Model_Num

Systems that support less than 3 separate 1P processors in a 2-processor socket system have the option of duplicating the one of the tested processors in multiple configuration fields in the QPL. The table below describes the options for the two possible scenarios.

	Low End	Typical	High End
Only 1 supported 1P Processor	A	A	A
Only 2 Supported 1P Processor	A	A or B	A or B